

Computational Implications of Reducing Data to Sufficient Statistics

Andrea Montanari*

July 31, 2015

Abstract

Given a large dataset and an estimation task, it is common to pre-process the data by reducing them to a set of sufficient statistics. This step is often regarded as straightforward and advantageous (in that it simplifies statistical analysis). I show that –on the contrary– reducing data to sufficient statistics can change a computationally tractable estimation problem into an intractable one. I discuss connections with recent work in theoretical computer science, and implications for some techniques to estimate graphical models.

1 Introduction

The idea of sufficient statistics is a cornerstone of statistical theory and statistical practice. Given a dataset, evaluating a set of sufficient statistics yields a concise representation that can be subsequently used to design (for instance) optimal statistical estimation procedures. To quote a widely adopted textbook [LC98]:

‘It often turns out that some part of the data carries no information about the unknown distribution and that \mathbf{X} can therefore be replaced by some statistic $\mathbf{T} = \mathbf{T}(\mathbf{X})$ without loss of information.’

The main point of the present paper is the following. While optimal statistical estimation can be performed solely on the basis of sufficient statistics, *reduction to sufficient statistics can lead to an explosion in computational complexity*. This phenomenon is so dramatic that –after reduction to sufficient statistics– a tractable estimation task can become entirely intractable.

To be concrete, we shall consider the problem of estimating the parameters of an exponential family over $\mathbf{x} \in \{0, 1\}^p$. Namely, we consider the following density with respect to the counting measure on $\{0, 1\}^p$:

$$p_{\boldsymbol{\theta}}(\mathbf{x}) = \frac{1}{Z(\boldsymbol{\theta})} h(\mathbf{x}) e^{\langle \boldsymbol{\theta}, \mathbf{x} \rangle}. \quad (1)$$

Here $h : \{0, 1\}^p \rightarrow \mathbb{R}_{>0}$ will be assumed strictly positive, $\boldsymbol{\theta} \in \mathbb{R}^p$ and $\langle \mathbf{a}, \mathbf{b} \rangle = \sum_{i=1}^p a_i b_i$ is the standard scalar product in \mathbb{R}^p . The partition function $Z(\boldsymbol{\theta})$ is determined by the normalization condition, yielding $Z(\boldsymbol{\theta}) = \sum_{\mathbf{x} \in \{0, 1\}^p} h(\mathbf{x}) e^{\langle \boldsymbol{\theta}, \mathbf{x} \rangle}$.

*Department of Electrical Engineering and Department of Statistics, Stanford University

We assume to be given n i.i.d. samples $\mathbf{X}_1, \dots, \mathbf{X}_n \sim p_{\boldsymbol{\theta}}$ and want to estimate $\boldsymbol{\theta}$. Introduced the notation $\mathbf{X}^{(n)} = (\mathbf{X}_1, \dots, \mathbf{X}_n)$ for the dataset of n samples, taking values in $\{0, 1\}^{p \times n}$. A consistent estimator is a function $\hat{\boldsymbol{\theta}} : \{0, 1\}^{p \times n} \rightarrow \mathbb{R}^p$ such that, almost surely with respect to $\mathbf{X}_1, \dots, \mathbf{X}_n \sim p_{\boldsymbol{\theta}}$, we have

$$\lim_{n \rightarrow \infty} \hat{\boldsymbol{\theta}}(\mathbf{X}^{(n)}) = \boldsymbol{\theta}. \quad (2)$$

A vector of sufficient statistics for the model $p_{\boldsymbol{\theta}}$ is clearly given by the empirical average

$$\mathbf{T}(\mathbf{X}^{(n)}) = \frac{1}{n} \sum_{\ell=1}^n \mathbf{X}_{\ell}. \quad (3)$$

Classical statistical theory shows that, for any estimator $\hat{\boldsymbol{\theta}}$ as above, we can construct a new estimator $\hat{\boldsymbol{\theta}}^{\text{new}}$ with the *same statistical properties*, that depends only on the sufficient statistics $\mathbf{T}(\mathbf{X}^{(n)})$. In particular, $\hat{\boldsymbol{\theta}}^{\text{new}} : [0, 1]^p \rightarrow \mathbb{R}^p$ and, if $\hat{\boldsymbol{\theta}}$ is consistent, then $\hat{\boldsymbol{\theta}}^{\text{new}}$ is also consistent. Namely, $p_{\boldsymbol{\theta}}$ -almost surely

$$\lim_{n \rightarrow \infty} \hat{\boldsymbol{\theta}}^{\text{new}}(\mathbf{T}(\mathbf{X}^{(n)})) = \boldsymbol{\theta}. \quad (4)$$

In this paper we are concerned with the computational complexity of evaluating such consistent estimators. In particular we will ask whether, for a given estimation task, there exist consistent estimators that can be computed in time polynomial in the model dimensions p . (We refer to Section 2 for formal definitions.)

Our main result is that there exist simple exponential families of the form (1) for which (under standard computational complexity assumptions) no polynomial-time estimator exists that uses only sufficient statistics. On the other hand, for the same models, polynomial time estimators exist that use the whole data $\mathbf{X}^{(n)}$.

In other words, the estimation problem is computationally tractable if we use the whole dataset, but becomes intractable once the data are reduced to sufficient statistics. This suggests that reducing data to sufficient statistics might be harmful.

1.1 Classical argument for sufficient statistics

It is instructive to review the standard argument used to prove that a reduction in complexity entails no loss of information [LC98], since this is already suggestive of a possible explosion in computational complexity. Given any estimator $\hat{\boldsymbol{\theta}}(\mathbf{X}^{(n)})$ of the parameters $\boldsymbol{\theta}$, the argument constructs a randomized estimator $\hat{\boldsymbol{\theta}}^{\text{new}}(\mathbf{T})$ that only depends on the sufficient statistics \mathbf{T} , and has the same distribution as $\hat{\boldsymbol{\theta}}(\mathbf{X}^{(n)})$, in particular the same risk. The construction is fairly simple. Given \mathbf{T} , we can sample $\widetilde{\mathbf{X}}^{(n)}$ (an independent copy of $\mathbf{X}^{(n)}$) from the law $p_{\boldsymbol{\theta}}(\cdot | \mathbf{T})$, conditional on the observed sufficient statistics $\mathbf{T}(\widetilde{\mathbf{X}}^{(n)}) = \mathbf{T}$. By definition of sufficient statistics, the conditional law $p_{\boldsymbol{\theta}}(\cdot | \mathbf{T})$ does not depend on $\boldsymbol{\theta}$, and hence we can sample $\widetilde{\mathbf{X}}^{(n)}$ without knowing $\boldsymbol{\theta}$. We then define $\hat{\boldsymbol{\theta}}^{\text{new}}(\mathbf{T}) \equiv \hat{\boldsymbol{\theta}}(\widetilde{\mathbf{X}}^{(n)})$. This has clearly the same risk as $\hat{\boldsymbol{\theta}}(\mathbf{X}^{(n)})$. In words, we were able to generate new data as informative as the original one using the sufficient statistics.

The problem with this argument is that sampling from the conditional distribution of $p_{\boldsymbol{\theta}}$ given $\mathbf{T}(\widetilde{\mathbf{X}}^{(n)}) = \mathbf{T}$ can be computationally hard. Indeed simple examples can be given for the weight

$h(\cdot)$ (see Section 3) that make sampling from $p_{\theta}(\cdot|\mathbf{T})$ –even approximately– impossible unless $\text{RP} = \text{NP}$ ($\text{RP} \neq \text{NP}$ this is a standard computational complexity assumption that is analogous to $\text{P} \neq \text{NP}$, and better suited for sampling problems, see Section 3 for more details). In other words, this argument is based on a reduction that is not computationally efficient. As we shall see, in some examples, no reduction is computationally efficient, not only the particular one given here.

1.2 Some examples

In order to build intuition on our main result, we present two examples of the model (1). Both fits in the general context of estimating discrete graphical models. This problem has been the object of considerable interest across statistics and machine learning during recent years [BMS08, BM09, RWL⁺10, ATHW12]. Our intractability results imply that, for both of these examples, one can construct families of parameters that cannot be estimated in polynomial time under standard computational assumptions.

Ising model (unknown magnetic fields). Let $\mathbf{J} = (J_{ij})_{i,j \in [p]}$ be a symmetric matrix $\mathbf{J} = \mathbf{J}^T$. For $x \in \{0, 1\}$, let $s(x) = 1 - 2x \in \{+1, -1\}$. The associated Ising model is the probability measure on $\{0, 1\}^p$ with density function:

$$p_{\mathbf{J}, \theta}(\mathbf{x}) = \frac{1}{\tilde{Z}(\mathbf{J}, \theta)} \exp \left\{ \sum_{1 \leq i < j \leq p} J_{ij} s(x_i) s(x_j) - \frac{1}{2} \sum_{i=1}^p \theta_i s(x_i) \right\}. \quad (5)$$

Ising models are one of the best studied models in statistical physics [Isi25, Gri06]. Since the seminal work of Hopfield [Hop82] and Hinton and Sejnowski [AHS85], they found application in numerous areas of machine learning, computer vision, clustering and spatial statistics.

Following the statistical physics terminology, the pairwise parameters J_{ij} are called ‘couplings’, and the singleton parameters θ_i are referred to as ‘magnetic fields.’ (The standard normalization for the latter does not include the factor $(-1/2)$, that is introduced here for convenience.) The probability distribution (5) can be written in the form (1) by letting

$$h(\mathbf{x}) = \exp \left\{ \sum_{1 \leq i < j \leq p} J_{ij} s(x_i) s(x_j) \right\}. \quad (6)$$

In other words, the problem of estimating the parameter θ can be described as estimating the magnetic fields in an Ising model with *known couplings*.

Special choices of the couplings \mathbf{J} have attracted most of the work within statistical physics and theoretical computer science. A possible construction is based on a graph $G = (V, E)$ with vertex set $V = [p] = \{1, \dots, p\}$. We then set

$$J_{ij} = \begin{cases} \beta & \text{if } (i, j) \in E, \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

For $\beta > 0$ this is the so-called *ferromagnetic Ising model*. For $\beta < 0$, we get the *anti-ferromagnetic Ising model*. In Section 3.1 we will show that estimating the parameters θ from sufficient statistics in the latter model is –in general– computationally intractable. As a consequence, the problem is intractable for the more general model (6).

Ising model (unknown couplings). A natural question –in view of the last example– is whether our frameworks sheds any light into the hardness of estimating the couplings of an Ising model. In order to clarify this point, consider the case of a zero-field model on variables $\tilde{\mathbf{s}} = (\tilde{s}_0, \tilde{s}_1, \tilde{s}_2, \dots, \tilde{s}_p) \in \{+1, -1\}^{p+1}$, with coupling matrix $\mathbf{J} = (J_{ij})_{i,j \in \{0,1,\dots,p\}}$ and no magnetic field:

$$q_{\mathbf{J}}(\tilde{\mathbf{s}}) = \frac{1}{Z(\mathbf{J})} \exp \left\{ \sum_{0 \leq i < j \leq p} J_{ij} \tilde{s}_i \tilde{s}_j \right\}. \quad (8)$$

Note that $-\tilde{\mathbf{s}}$ is distributed as $\tilde{\mathbf{s}}$. Hence, without loss of generality, we can preprocess the data flipping them arbitrarily. Equivalently, we can consider the distribution of the vector $\mathbf{s} = (s_1, s_2, \dots, s_p)$, whereby $s_i = \tilde{s}_0 \tilde{s}_i$. If $\tilde{\mathbf{s}}$ has distribution $q_{\mathbf{J}}$, then \mathbf{s} has distribution

$$p_{\mathbf{J}}(\mathbf{s}) = \frac{1}{Z(\mathbf{J})} \exp \left\{ \sum_{1 \leq i < j \leq p} J_{ij} s_i s_j + \sum_{i=1}^p J_{0i} s_i \right\}. \quad (9)$$

Let us now assume that all the couplings have been estimated except the $(J_{0i})_{i \in [p]}$. Accordingly, we define the parameters $\theta_i = -2J_{0i}$ and rewrite the above distribution as

$$p_{\mathbf{J}_{\setminus 0}, \boldsymbol{\theta}}(\mathbf{s}) = \frac{1}{Z(\mathbf{J}_{\setminus 0}, \boldsymbol{\theta})} \exp \left\{ \sum_{1 \leq i < j \leq p} J_{ij} s_i s_j - \frac{1}{2} \sum_{i=1}^p \theta_i s_i \right\}, \quad (10)$$

where $\mathbf{J}_{\setminus 0} \equiv (J_{ij})_{i,j \in [p]}$. This coincides with Eq. (5) apart from the trivial change of variables $s_i = s(x_i)$. It follows that estimating a row of the coupling matrix, with knowledge of the other couplings, in absence of singleton terms (magnetic fields) is equivalent to estimating the singleton terms knowing all the couplings (in a model with one less variable).

The construction in Section 3.1 shows that there are choices of the couplings $\mathbf{J}_{\setminus 0}$ for which estimating the parameters $\boldsymbol{\theta}$ from sufficient statistics is intractable. Equivalently, it is (in general) intractable to estimate $(J_{0i})_{i \in [p]}$ using the knowledge of $(J_{ij})_{i,j \in [p]}$, and of the empirical averages of $(s_i)_{i \in [p]} = (\tilde{s}_0 \tilde{s}_i)_{i \in [p]}$.

1.3 Paper outline

The rest of the paper is organized as follows. In Section 2 we state formally our main results. Under technical assumptions this shows that, if there is a computationally efficient estimator $\hat{\boldsymbol{\theta}}(\mathbf{T})$ that use only sufficient statistics, then the partition function $Z(\mathbf{0})$ can be approximated, also efficiently. In Section 3 we construct a family of functions $h(\cdot)$ for which approximating $Z(\mathbf{0})$ in polynomial time is impossible unless $\text{RP} = \text{NP}$. As a consequence of our main theorem, no efficient parameter estimator from sufficient statistics exists for these models (unless $\text{RP} = \text{NP}$). Remarkably, we show that a simple consistent estimator can be developed using the data $\mathbf{X}^{(n)}$, for the same models.

The example in Section 3 is an undirected graphical model, and indeed intractability of computing approximations to $Z(\mathbf{0})$ is quite generic in this context. In Section 4 we discuss relations to the literature in algorithms and graphical models estimation.

1.4 Notations

We will use $[n] = \{1, 2, \dots, n\}$ to denote the set of first n integers. Whenever possible, we will follow the convention of denoting deterministic values by lowercase letters (e.g. x, y, z, \dots) and random

variables by uppercase letters (e.g. X, Y, Z, \dots). We reserve boldface for vectors and matrices, e.g. \mathbf{x} is a deterministic vector or matrix, and \mathbf{X} a random vector or matrix. For a vector $\mathbf{x} \in \mathbb{R}^m$, and a set $A \subseteq [m]$, \mathbf{x}_A denotes the subvector indexed by A . The components of \mathbf{x} are denoted by (x_1, x_2, \dots, x_m) .

Given $f : \mathbb{R}^m \rightarrow \mathbb{R}$, we denote by $\nabla f(\mathbf{x}) \in \mathbb{R}^m$ its gradient at \mathbf{x} , and by $\nabla^2 f(\mathbf{x}) \in \mathbb{R}^{m \times m}$ its Hessian. Whenever useful, we add as a subscript the variable with respect to which we are differentiating as, for instance, in $\nabla_{\mathbf{x}} f(\mathbf{x})$.

2 Computational hardness of consistent parameter estimation

As mentioned above, we consider the problem of consistent estimation of the parameters $\boldsymbol{\theta}$ of the model (1), from the sufficient statistics (3). Throughout we will assume $h(\mathbf{x}) > 0$ strictly for all $\mathbf{x} \in \{0, 1\}^p$.

As $n \rightarrow \infty$, $\mathbf{T}(\mathbf{X}^{(n)})$ converges –by the law of large numbers– to the population mean

$$\boldsymbol{\tau}_*(\boldsymbol{\theta}) \equiv \mathbb{E}_{\boldsymbol{\theta}}\{\mathbf{X}\} = \frac{1}{Z(\boldsymbol{\theta})} \sum_{\mathbf{x} \in \{0, 1\}^p} \mathbf{x} h(\mathbf{x}) e^{\langle \boldsymbol{\theta}, \mathbf{x} \rangle}. \quad (11)$$

Here and below $\mathbb{E}_{\boldsymbol{\theta}}$ denotes expectation with respect to the distribution $p_{\boldsymbol{\theta}}$. We use the $*$ subscript to emphasize that this $\boldsymbol{\tau}_*$ is a function, and distinguish it from a specific value $\boldsymbol{\tau} \in (0, 1)^p$. A consistent estimator based on the sufficient statistics \mathbf{T} must necessarily invert this function in the limit $n \rightarrow \infty$. This motivates the following definition.

Definition 2.1. For $p \in \mathbb{N}$, let $\mathcal{H}_p = \{h(\cdot)\}$ be a set of functions $h : \{0, 1\}^p \rightarrow \mathbb{R}_{>0}$. A polynomial consistent estimator from sufficient statistics for the model (1) with $h \in \mathcal{H}_p$ is an algorithm that given $\boldsymbol{\tau} = (\tau_1, \tau_2, \dots, \tau_p) \in (0, 1)^p$, and a precision parameter $\xi > 0$, returns $\hat{\boldsymbol{\theta}}(\boldsymbol{\tau}, \xi)$ such that

1. $\|\hat{\boldsymbol{\theta}}(\boldsymbol{\tau}, \xi) - \boldsymbol{\theta}\|_2 \leq \xi$ for some $\boldsymbol{\theta}$ such that $\boldsymbol{\tau}_*(\boldsymbol{\theta}) = \boldsymbol{\tau}$.
2. The algorithm terminates in time polynomial in $1/\xi$, p , the maximum description length of any of the τ_i , and the description length of $h \in \mathcal{H}_p$.

Two remarks are in order. The terminology adopted here (in particular, the use of ‘consistent’ and ‘polynomial’) is motivated by the following remarks. First, in the terminology of complexity theory, this is a ‘fully polynomial time approximation scheme.’

Second, as discussed below, there is indeed a unique, continuous, function $\boldsymbol{\theta}_* : (0, 1)^p \rightarrow \mathbb{R}^p$ such that $\boldsymbol{\tau}_*(\boldsymbol{\theta}_*(\boldsymbol{\tau})) = \boldsymbol{\tau}$. This implies that $\boldsymbol{\theta}$ is indeed a consistent estimator in the sense that for a sequence $\xi_n \rightarrow 0$, we have, almost surely and hence in probability,

$$\lim_{n \rightarrow \infty} \hat{\boldsymbol{\theta}}(\mathbf{T}(\mathbf{X}^{(n)}), \xi_n) = \boldsymbol{\theta}. \quad (12)$$

Finally, in the following we shall occasionally drop the qualification ‘from sufficient statistics’ when it is clear that we are considering estimators that only use sufficient statistics.

We next state our main theorem.

Theorem 1. Assume \mathcal{H}_p , $p \geq 1$, to be a family of weight functions such that the the following three conditions hold:

C1. There exists $\delta \in (0, 1/2)$ such that, for any $p \in \mathbb{N}$, any $h \in \mathcal{H}_p$, and any $i \in [p]$

$$\tau_{*,i}(\boldsymbol{\theta} = \mathbf{0}) \in (\delta, 1 - \delta). \quad (13)$$

C2. There exists a polynomial $L(p)$ such that, for any $p \in \mathbb{N}$, any $h \in \mathcal{H}_p$, we have, for all $\boldsymbol{\theta}$ such that $\boldsymbol{\tau}_*(\boldsymbol{\theta}) \in [\delta, 1 - \delta]^p$,

$$\text{Cov}_{\boldsymbol{\theta}}(\mathbf{X}) \equiv \mathbb{E}_{\boldsymbol{\theta}}\{\mathbf{X}\mathbf{X}^\top\} - \mathbb{E}_{\boldsymbol{\theta}}\{\mathbf{X}\}\mathbb{E}_{\boldsymbol{\theta}}\{\mathbf{X}\}^\top \succeq \frac{1}{L(p)} \mathbf{I}_{p \times p}. \quad (14)$$

C3. There exists a polynomial $K(p)$ such that, for any $p \in \mathbb{N}$, any $h \in \mathcal{H}_p$

$$\max_{\mathbf{x}, \mathbf{y} \in \{0,1\}^p} |\log h(\mathbf{x}) - \log h(\mathbf{y})| \leq K(p). \quad (15)$$

If there exists a polynomial consistent estimator for the model $\mathcal{H} = (\mathcal{H}_p)_{p \geq 1}$, then, for any $\varepsilon > 4p\delta(K(p) + \log(4/\delta))$, there exists an algorithm polynomial in the description length of h , and in $(1/\varepsilon)$, returning \hat{Z} such that

$$e^{-\varepsilon} Z(\mathbf{0}) \leq \hat{Z} \leq e^\varepsilon Z(\mathbf{0}). \quad (16)$$

The main implication of this theorem is negative. If the problem of approximating $Z(\mathbf{0})$ is intractable, it follows from the above that there is no polynomial consistent estimator. We will show in Section 3 that we can use recent results in complexity theory to construct a fairly simple class $\mathcal{H} = \{h\}$ such that:

- An approximation scheme does not exist for $Z(\mathbf{0})$ under a standard complexity theory assumption known as $\text{RP} \neq \text{NP}$ (and analogous to $\text{P} \neq \text{NP}$, see Section 3).
- Remarkably, a consistent and computationally efficient estimator $\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}(\mathbf{X}^{(n)})$ of the model parameters exists! However, this is not a function only of the sufficient statistics.

As a direct consequence, we have the following.

Corollary 2.2. *There are classes of models \mathcal{H} for which no polynomial consistent estimator of the parameters from sufficient statistics exists, unless $\text{RP} = \text{NP}$.*

Remark 2.3. Note that, neither the assumption nor the conclusion of Theorem 1 (or of Corollary 2.2) depend on the number of samples n . In particular, the map $\boldsymbol{\tau} \mapsto \hat{\boldsymbol{\theta}}(\boldsymbol{\tau}, \xi)$ in Definition 2.1 is a map that approximately inverts the moment map $\boldsymbol{\theta} \mapsto \boldsymbol{\tau}_*(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta}}\{\mathbf{X}\}$.

Therefore, Theorem 1 and Corollary 2.2 can be read as stating that the map $\boldsymbol{\theta} \mapsto \boldsymbol{\tau}_*(\boldsymbol{\theta})$ cannot be (approximately) inverted in polynomial time.

The rest of this section is devoted to the proof of Theorem 1. It is useful to first recall a few well-known properties of exponential families, as specialized to the present setting. While these are consequences of fairly standard general statements (see for instance [Efr78, LC98]), we present self-contained proofs in Appendix A for the readers' convenience.

We will use standard statistics notations for the *cumulant generating function* (also known as *log-partition function*), $A : \mathbb{R}^p \rightarrow \mathbb{R}$,

$$A(\boldsymbol{\theta}) \equiv \log Z(\boldsymbol{\theta}) = \log \left\{ \sum_{\mathbf{x} \in \{0,1\}^p} h(\mathbf{x}) e^{\langle \boldsymbol{\theta}, \mathbf{x} \rangle} \right\}, \quad (17)$$

and its Legendre-Fenchel transform $F : (0, 1)^p \rightarrow \mathbb{R}$, which we shall call the *free energy*,

$$F(\boldsymbol{\tau}) = \inf_{\boldsymbol{\theta} \in \mathbb{R}^p} (A(\boldsymbol{\theta}) - \langle \boldsymbol{\tau}, \boldsymbol{\theta} \rangle). \quad (18)$$

Proposition 2.4. *Given a strictly positive $h : \{0, 1\}^p \rightarrow \mathbb{R}_{>0}$, the following hold:*

Fact1. *The function $\boldsymbol{\tau}_* : \mathbb{R}^p \rightarrow (0, 1)^p$ is $C_\infty(\mathbb{R}^p)$.*

Fact2. *Recalling that $\text{Cov}_{\boldsymbol{\theta}}(\mathbf{X}) \in \mathbb{R}^{p \times p}$ denotes the covariance matrix of the law $p_{\boldsymbol{\theta}}$, we have*

$$\nabla_{\boldsymbol{\theta}} A(\boldsymbol{\theta}) = \boldsymbol{\tau}_*(\boldsymbol{\theta}), \quad (19)$$

$$\nabla_{\boldsymbol{\theta}}^2 A(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} \boldsymbol{\tau}_*(\boldsymbol{\theta}) = \text{Cov}_{\boldsymbol{\theta}}(\mathbf{X}). \quad (20)$$

Fact3. *$\text{Cov}_{\boldsymbol{\theta}}(\mathbf{X}) \succeq c(\boldsymbol{\theta}) \mathbf{I}_p$ for some continuous strictly positive $c(\boldsymbol{\theta}) > 0$.*

Fact4. *$\boldsymbol{\tau}_* : \mathbb{R}^p \rightarrow (0, 1)^p$ is a bijection. We will denote by $\boldsymbol{\theta}_* : (0, 1)^p \rightarrow \mathbb{R}^p$ the inverse mapping.*

Fact5. *The function $F : (0, 1)^p \rightarrow \mathbb{R}$ is concave and $C_\infty((0, 1)^p)$ with*

$$A(\boldsymbol{\theta}) = \max_{\boldsymbol{\tau} \in (0, 1)^p} \{F(\boldsymbol{\tau}) + \langle \boldsymbol{\tau}, \boldsymbol{\theta} \rangle\} = F(\boldsymbol{\tau}_*(\boldsymbol{\theta})) + \langle \boldsymbol{\tau}_*(\boldsymbol{\theta}), \boldsymbol{\theta} \rangle, \quad (21)$$

$$\nabla_{\boldsymbol{\tau}} F(\boldsymbol{\tau}) = -\boldsymbol{\theta}_*(\boldsymbol{\tau}) \quad (22)$$

$$\nabla_{\boldsymbol{\tau}}^2 F(\boldsymbol{\tau}) = -\nabla_{\boldsymbol{\tau}} \boldsymbol{\theta}_*(\boldsymbol{\tau}) = -[\nabla_{\boldsymbol{\theta}}^2 A(\boldsymbol{\theta}_*(\boldsymbol{\tau}))]^{-1}. \quad (23)$$

Fact6. *Let $\mathcal{P}_p = \mathcal{P}(\{0, 1\}^p)$ be the set of probability distributions over $\{0, 1\}^p$, and denote by $H(q) = -\sum_{\mathbf{x} \in \{0, 1\}^p} q(\mathbf{x}) \log q(\mathbf{x})$ the Shannon entropy of $q \in \mathcal{P}_p$. Then*

$$F(\boldsymbol{\tau}) = \max_{q \in \mathcal{P}_p} \left\{ H(q) + \mathbb{E}_q \log h(\mathbf{X}) \text{ such that } \mathbb{E}_q \{\mathbf{X}\} = \boldsymbol{\tau} \right\}. \quad (24)$$

We are now in position to prove Theorem 1: we present here a version of the proof that omits some technical step, and complete the details in Appendix B.

Proof of Theorem 1. By Eq. (21) we have, letting $D_p(\delta) \equiv [\delta, 1 - \delta]^p$,

$$\log Z(\mathbf{0}) = \max_{\boldsymbol{\tau} \in (0, 1)^p} F(\boldsymbol{\tau}) \quad (25)$$

$$= \max_{\boldsymbol{\tau} \in D_p(\delta)} F(\boldsymbol{\tau}) = F(\boldsymbol{\tau}_*(\mathbf{0})), \quad (26)$$

where the second equality follows from assumption C1.

The problem of computing $Z(\mathbf{0})$ is then reduced to the one of maximizing the concave function $F(\boldsymbol{\tau})$ over the convex set $D_p(\delta)$. Note that, by assumption C2 and Eq. (23), the gradient of $F(\boldsymbol{\tau})$ has Lipschitz modulus bounded by $L(p)$ on $D_p(\delta)$. Namely, for all $\boldsymbol{\tau}_1, \boldsymbol{\tau}_2 \in D_p(\delta)$,

$$\|\nabla_{\boldsymbol{\tau}} F(\boldsymbol{\tau}_1) - \nabla_{\boldsymbol{\tau}} F(\boldsymbol{\tau}_2)\|_2 \leq L(p) \|\boldsymbol{\tau}_1 - \boldsymbol{\tau}_2\|_2. \quad (27)$$

We will maximize $F(\boldsymbol{\tau})$ by a standard projected gradient algorithm. We will work here under the assumption that we have access to an oracle that given a point $\boldsymbol{\tau} \in D_p(\delta)$, returns $\boldsymbol{\theta}_*(\boldsymbol{\tau}) =$

$-\nabla_{\boldsymbol{\tau}} F(\boldsymbol{\tau})$. While in reality we do not have access to such an oracle, Definition 2.1 (and the Theorems assumptions) imply that there exists an efficient approximation scheme for this oracle. We will show in Appendix B that indeed we can replace the oracle by such an approximation scheme.

Given the oracle, the projected gradient algorithm is defined by the iteration (with the superscript t indicating the iteration number, and letting $L = L(p)$)

$$\boldsymbol{\tau}^0 = \left(\frac{1}{2}, \frac{1}{2}, \dots, \frac{1}{2}\right)^\top, \quad (28)$$

$$\boldsymbol{\tau}^{t+1} = \mathbf{P}_\delta \left(\boldsymbol{\tau}^t - \frac{1}{L} \boldsymbol{\theta}_*(\boldsymbol{\tau}^t) \right). \quad (29)$$

Here $\mathbf{P}_\delta(\mathbf{x})$ is the orthogonal projector on $D_p(\delta)$, which can be computed efficiently since, for each $i \in \{1, 2, \dots, p\}$, and $\mathbf{u} \in \mathbb{R}^p$,

$$\mathbf{P}_\delta(\mathbf{u})_i = \begin{cases} \delta & \text{if } u_i < \delta, \\ u_i & \text{if } u_i \in [\delta, 1 - \delta], \\ 1 - \delta & \text{if } u_i > 1 - \delta. \end{cases} \quad (30)$$

We run $t_0 = t_0(p, \varepsilon) \equiv \lceil 2pL(p)/\varepsilon \rceil$ iterations of projected gradient. By [BT09, Theorem 3.1] we have

$$0 \leq F(\boldsymbol{\tau}_*(\mathbf{0})) - F(\boldsymbol{\tau}^{t_0}) \leq \frac{L \|\boldsymbol{\tau}^0 - \boldsymbol{\tau}_*(\mathbf{0})\|_2^2}{2t_0} \leq \frac{\varepsilon}{4}. \quad (31)$$

Hence, $F(\boldsymbol{\tau}^{t_0})$ provides a good approximation of $F(\boldsymbol{\tau}_*(\mathbf{0})) = \log Z(\mathbf{0})$.

We are left with the task of evaluating $F(\boldsymbol{\tau}^{t_0})$. The idea is to ‘integrate’ the derivative of $F(\boldsymbol{\tau})$ along a path that starts at a point $\boldsymbol{\tau}^{(0)}$ where F can be easily approximated. We will use again $\nabla_{\boldsymbol{\tau}} F(\boldsymbol{\tau}) = -\boldsymbol{\theta}_*(\boldsymbol{\tau})$ and assume that $\boldsymbol{\theta}_*$ is exactly computed by an oracle. Again we will see in Appendix B that this oracle can be replaced by the estimator in Definition 2.1

Let $\boldsymbol{\tau}^{(0)} \equiv (\delta, \dots, \delta)^\top$. Next consider Eq. (24). For any $q \in \mathcal{P}_p$ such that $\mathbb{E}_q(\mathbf{X}) = \boldsymbol{\tau}^{(0)}$, we have, letting $s(x) \equiv -x \log x - (1-x) \log(1-x)$

$$0 \leq H(q) \leq p s(\delta), \quad (32)$$

$$\left| \mathbb{E}_p \log \frac{h(\mathbf{X})}{h(\mathbf{0})} \right| \leq K(p) \mathbb{P}_q(\mathbf{X} \neq \mathbf{0}) \leq K(p) p \delta. \quad (33)$$

Where the second inequality in Eq. (32) follows since the joint entropy is not larger than the sum of the entropy of the marginals. The first inequality in Eq. (33) follows from assumption C3, and the last inequality in Eq. (33) is Markov’s. Using these bounds in Eq. (24), we get

$$\left| F(\boldsymbol{\tau}^{(0)}) - \log h(\mathbf{0}) \right| \leq p s(\delta) + p K(p) \delta \leq \frac{\varepsilon}{4}, \quad (34)$$

where the last inequality follows from the assumption $\varepsilon > 4p\delta(K(p) + \log(4/\delta))$.

For an integer m , we let $\boldsymbol{\tau}^{(m)} = \boldsymbol{\tau}^{t_0}$ be the output of projected gradient, and, for $\ell \in \{1, \dots, m-1\}$, $\boldsymbol{\tau}^{(\ell)} \equiv \boldsymbol{\tau}^{(0)} + \ell(\boldsymbol{\tau}^{(m)} - \boldsymbol{\tau}^{(0)})/m$ be given by linearly interpolating between $\boldsymbol{\tau}^{(0)}$ and $\boldsymbol{\tau}^{(m)}$. Note that, since $\boldsymbol{\tau}^{(0)}, \boldsymbol{\tau}^{(m)} \in D_p(\delta)$, we have $\boldsymbol{\tau}^{(\ell)} \in D_p(\delta)$ as well, by convexity.

We finally construct our approximation \widehat{Z} by letting

$$\log \widehat{Z}^{\text{or}} \equiv \log h(\mathbf{0}) - \sum_{\ell=1}^m \langle \boldsymbol{\theta}_*(\boldsymbol{\tau}^{(\ell)}), \boldsymbol{\tau}^{(\ell)} - \boldsymbol{\tau}^{(\ell-1)} \rangle. \quad (35)$$

(We introduced the superscript or to emphasize that this approximation makes use of the oracle $\boldsymbol{\theta}_*$. In Appendix B we control the additional error induced by the use of $\hat{\boldsymbol{\theta}}$.) Let us bound the approximation error:

$$\begin{aligned} \left| \log \frac{\widehat{Z}^{\text{or}}}{Z(\mathbf{0})} \right| &\leq |F(\boldsymbol{\tau}^{(0)}) - \log h(\mathbf{0})| \\ &\quad + \sum_{\ell=1}^m |F(\boldsymbol{\tau}^{(\ell)}) - F(\boldsymbol{\tau}^{(\ell-1)}) - \langle \nabla_{\boldsymbol{\tau}} F(\boldsymbol{\tau}^{(\ell)}), \boldsymbol{\tau}^{(\ell)} - \boldsymbol{\tau}^{(\ell-1)} \rangle| \\ &\quad + |F(\boldsymbol{\tau}_*(\mathbf{0})) - F(\boldsymbol{\tau}^{(m)})|. \end{aligned} \quad (36)$$

The first term is bounded by Eq. (34) and the last by Eq. (31). As for the middle term, by the intermediate value theorem there exists, for each ℓ , a point $\tilde{\boldsymbol{\tau}}^{(\ell)} \in [\boldsymbol{\tau}^{(\ell-1)}, \boldsymbol{\tau}^{(\ell)}]$ such that

$$\begin{aligned} |F(\boldsymbol{\tau}^{(\ell)}) - F(\boldsymbol{\tau}^{(\ell-1)}) - \langle \nabla_{\boldsymbol{\tau}} F(\boldsymbol{\tau}^{(\ell)}), \boldsymbol{\tau}^{(\ell)} - \boldsymbol{\tau}^{(\ell-1)} \rangle| &= |\langle \nabla_{\boldsymbol{\tau}} F(\tilde{\boldsymbol{\tau}}^{(\ell)}) - \nabla_{\boldsymbol{\tau}} F(\boldsymbol{\tau}^{(\ell)}), \boldsymbol{\tau}^{(\ell)} - \boldsymbol{\tau}^{(\ell-1)} \rangle| \\ &\leq L(p) \|\boldsymbol{\tau}^{(\ell)} - \boldsymbol{\tau}^{(\ell-1)}\|_2^2 \leq \frac{L(p)}{m^2} \|\boldsymbol{\tau}^{(m)} - \boldsymbol{\tau}^{(0)}\|_2^2 \leq \frac{L(p)p}{m^2}. \end{aligned} \quad (37)$$

Hence, choosing $m = m_0(p, \varepsilon) \equiv \lceil 4L(p)p/\varepsilon \rceil$, we can bound the sum in Eq. (36) by $\varepsilon/4$. Hence the approximation error is bounded by

$$\left| \log \frac{\widehat{Z}^{\text{or}}}{Z(\mathbf{0})} \right| \leq \frac{\varepsilon}{4} + \frac{\varepsilon}{4} + \frac{\varepsilon}{4} \leq \frac{3\varepsilon}{4}, \quad (38)$$

which concludes our proof outline. \square

3 An example

In this section we describe a simple class of functions \mathcal{H}_p for which it is impossible to estimate $Z(\mathbf{0})$ in polynomial time unless $\text{RP} = \text{NP}$. $\text{RP} \neq \text{NP}$ is a standard computational assumption analogous to $\text{P} \neq \text{NP}$. Informally, NP is the set of decision problems whose solution can be checked in polynomial time. On the other hand, RP is the set of problems that can be solved in polynomial time using a randomized algorithm (roughly, an algorithm that has access to an ideal random numbers generator). The assumption $\text{RP} \neq \text{NP}$ posits that there are problems that cannot be solved in polynomial time even using randomization. As such, it is a slightly stronger than $\text{P} \neq \text{NP}$ but nearly as widely believed to be true. We refer to standard textbooks, e.g. [AB09], for further details.

Using our Theorem 1, we will show that consistent parameter estimation using sufficient statistics is intractable for the models in the class \mathcal{H}_p . We will then show that –for the same models– it is quite easy to estimate the parameters from i.i.d. samples $\mathbf{X}_1, \dots, \mathbf{X}_n \sim p_{\boldsymbol{\theta}}$.

3.1 Intractability of estimation from sufficient statistics

For $\beta > 0$ a fixed known number, and $G = (V = [p], E)$ a simple graph, we let $h_{G,\beta} : \{0, 1\}^p \rightarrow \mathbb{R}_{>0}$ be defined by

$$h_{G,\beta}(\mathbf{x}) \equiv \exp \left\{ 2\beta \sum_{(i,j) \in E} \mathbb{I}(x_i \neq x_j) \right\}. \quad (39)$$

This is also known as the *anti-ferromagnetic Ising model* on graph G . (The same model was already introduced in Section 1.2, with slightly different normalizations.)

Fixing $k \in \mathbb{N}$, and $\beta \in \mathbb{R}_{>0}$, we introduce the class of functions

$$\mathcal{H}_p(k, \beta) \equiv \{h_{G,\beta} : G \text{ is a simple regular graph of degree } k \text{ over } p \text{ vertices}\}. \quad (40)$$

(Recall that a regular graph is a graph with the same degree at all vertices. The set of regular graphs is non-empty as soon as pk is even, and $p \geq p_0(k)$.) Intractability of approximating $Z(\mathbf{0})$ was characterized in [SS12, GSV12]. We restate the main result of [SS12], adapting it to the present setting¹.

Theorem 2 (Sly, Sun, 2012). *For any $k \geq 3$ and $\beta > \operatorname{atanh}(1/(k-1))$ there exists $\varepsilon_0 = \varepsilon_0(k, \beta) > 0$ such that the following holds. Unless $\text{RP} = \text{NP}$, there is no polynomial algorithm taking as input $h \in \mathcal{H}_p(k, \beta)$ and returning \hat{Z} such that*

$$e^{-\varepsilon_0 p} Z(\mathbf{0}) \leq \hat{Z} \leq e^{\varepsilon_0 p} Z(\mathbf{0}) \quad (41)$$

where we recall that $Z(\mathbf{0}) \equiv \sum_{\mathbf{x} \in \{0,1\}^p} h(\mathbf{x})$.

We note in passing that [SST14] implies that this result is tight. Namely, for $\beta < \operatorname{atanh}(1/(k-1))$ there exists a fully polynomial approximation scheme for $Z(\mathbf{0})$.

We can now state (and prove) a more concrete form of Corollary 2.2. In Section 3.2 we will show that the model parameters can be consistently estimated in polynomial time from i.i.d. samples $\mathbf{X}_1, \dots, \mathbf{X}_n \sim p_{\boldsymbol{\theta}}$.

Corollary 3.1. *Fix $k \geq 3$ and $\beta > \operatorname{atanh}(1/(k-1))$. Unless $\text{RP} = \text{NP}$, then there exists no polynomial consistent estimator from sufficient statistics for the model $\{\mathcal{H}_p(k, \beta)\}_{p \geq p_0(k)}$.*

Proof. The proof consists in checking that the assumptions C1, C2, C3 of Theorem 1 apply. It then follows from Theorem 2 that either $\text{RP} = \text{NP}$ or there is no polynomial consistent estimator from sufficient statistics. Throughout, we will write C or c for generic strictly positive constants that can depend on β, k .

Let us start with condition C1. We fix a graph $G = (V = [p], E)$ on p vertices. For a vertex $i \in [p]$, we let $\partial i = \{j \in [p] : (i, j) \in E\}$ be the set of neighbors of i . Writing $\mathbb{P}_{\mathbf{0}}$ for $\mathbb{P}_{\boldsymbol{\theta}=\mathbf{0}}$, we have $\tau_{*,i}(\mathbf{0}) = \mathbb{P}_{\mathbf{0}}(X_i = 1)$.

¹The present statement differs from the original one in [SS12] in one technical aspect. We state that $Z(\mathbf{0})$ cannot be approximated within a ratio $e^{\varepsilon_0 p}$, while [SS12] state their result in terms of FPRAS. However, the authors of [SS12] also mention that their same proof yields impossibility to achieve the approximation ratio stated here (which is in fact easy to check).

Note that p_{θ} is a Markov Random Field with graph G . We therefore have

$$\min_{\mathbf{x}_{\partial i} \in \{0,1\}^{\partial i}} \mathbb{P}_{\mathbf{0}}(X_i = 1 | \mathbf{X}_{\partial i} = \mathbf{x}_{\partial i}) \leq \mathbb{P}_{\mathbf{0}}(X_i = 1) \leq \max_{\mathbf{x}_{\partial i} \in \{0,1\}^{\partial i}} \mathbb{P}_{\mathbf{0}}(X_i = 1 | \mathbf{X}_{\partial i} = \mathbf{x}_{\partial i}). \quad (42)$$

A simple direct calculation with Eq. (39) yields

$$\mathbb{P}_{\mathbf{0}}(X_i = 1 | \mathbf{X}_{\partial i} = \mathbf{x}_{\partial i}) = \frac{e^{2\beta n_0(\mathbf{x}_{\partial i})}}{e^{2\beta n_0(\mathbf{x}_{\partial i})} + e^{2\beta n_1(\mathbf{x}_{\partial i})}}, \quad (43)$$

where $n_0(\mathbf{x}_{\partial i})$ and $n_1(\mathbf{x}_{\partial i})$ are the number of zeros and ones in the vector $\mathbf{x}_{\partial i}$.

The right hand side of Eq. (43) is maximized when $n_0(\mathbf{x}_{\partial i}) = k - n_1(\mathbf{x}_{\partial i}) = k$, and minimized when $n_0(\mathbf{x}_{\partial i}) = k - n_1(\mathbf{x}_{\partial i}) = 0$. We then have

$$\frac{1}{1 + e^{2\beta k}} \leq \mathbb{P}_{\mathbf{0}}(X_i = 1) \leq \frac{e^{2\beta k}}{1 + e^{2\beta k}}. \quad (44)$$

We conclude that there exists $c = c(k, \beta) > 0$ such that condition C1 is satisfied for all $\delta \in (0, c)$. We will select the value of δ after checking condition C3.

Consider condition C3. From Eq. (39), it follows that $h_{G,\beta}(\mathbf{x}) \geq 1$ (because the argument in the exponent is non-negative) and $h_{G,\beta}(\mathbf{x}) \leq \exp(2\beta|E|)$ with $|E|$ the number of edges (the upper bound is saturated if the graph is bipartite). Since $|E| = kp/2$, Eq. (15) follows with $K(p) = \beta kp$.

At this point we choose $\delta = 1/(10pK(p))$. Assuming, without loss of generality, $K(p), p \geq 10$, this implies that we can take $\varepsilon = 2$ in Eq. (16), which yields the desired contradiction with Theorem 2, provided we can verify condition C2 with the stated value of δ .

To conclude our proof, consider condition C2. First we claim that $\tau_*(\theta) \in [\delta, 1 - \delta]^p$ implies $\|\theta\|_{\infty} \leq C \log p$ for some constant $C = C(k, \beta)$. Indeed, let us fix $i \in [p]$ and prove that $\theta_i \leq C \log p$ (the lower bound follows from an analogous argument). Using again the fact that p_{θ} is a Markov Random Field, we have (since, by definition $\tau_{*,i}(\theta) = \mathbb{P}_{\theta}(X_i = 1)$)

$$1 - \delta \geq \mathbb{P}_{\theta}(X_i = 1) \geq \min_{\mathbf{x}_{\partial i} \in \{0,1\}^{\partial i}} \mathbb{P}_{\theta}(X_i = 1 | \mathbf{X}_{\partial i} = \mathbf{x}_{\partial i}). \quad (45)$$

Proceeding as in the case of condition C1, we see that the last conditional probability is minimized when all the neighbors of i are in state 0 (i.e. $\mathbf{x}_{\partial i} = \mathbf{0}$). This yields

$$1 - \delta \geq \frac{e^{\theta_i}}{e^{\theta_i} + e^{2\beta k}}, \quad (46)$$

which is equivalent to $\theta_i \leq 2\beta k + \log((1 - \delta)/\delta)$. Substituting $\delta = 1/(10pK(p))$ with $K(p)$ a polynomial yields the desired bound (recall that β and k are constants).

Next, we claim that, there exists a polynomial $L_0(p)$ such that, for each $i \in [p]$, all $\mathbf{x}_{\partial i}$ and all θ with $|\theta_i| \leq C \log p$, we have

$$\text{Var}_{\theta}(X_i | \mathbf{X}_{\partial i} = \mathbf{x}_{\partial i}) \geq \frac{1}{L_0(p)}. \quad (47)$$

The calculation is essentially the same as the one already carried out above for $\mathbb{P}_{\theta}(X_i = 1 | \mathbf{X}_{\partial i} = \mathbf{x}_{\partial i})$ and therefore we omit it.

Finally, we want to prove Eq. (14) for some polynomial $L(p)$. Equivalently, we need to prove that $\text{Var}_\theta(\langle \mathbf{v}, \mathbf{X} \rangle) \geq 1/L(p)$ for any vector $\mathbf{v} \in \mathbb{R}^p$, $\|\mathbf{v}\|_2 = 1$. Fix one such vector \mathbf{v} . Let $i(1)$ be the index of the component of \mathbf{v} with largest magnitude (i.e. $|v_{i(1)}| \geq |v_j|$ for all $j \in [p] \setminus \{i(1)\}$). Let $i(2)$ be the index of the component of \mathbf{v} with largest magnitude *excluded* $i(1)$ and $\partial i(1)$ (i.e. $|v_{i(2)}| \geq |v_j|$ for all $j \in [p] \setminus \{i(1), \partial i(1), i(2)\}$), and so on. Namely, for each ℓ , we let $i(\ell)$ be the index of the component of \mathbf{v} with the largest magnitude *excluded* $i(1), \dots, i(\ell-1)$ and *their neighbors in G* . Denote by $m \geq n/(k+1)$ the total number of vertices selected in this manner, and let $S = \{i(1), \dots, i(m)\}$. It is immediate to see that

$$\sum_{i \in S} v_i^2 \geq \frac{\|\mathbf{v}\|_2^2}{k+1}. \quad (48)$$

Further, letting $S^c = [p] \setminus S$,

$$\text{Var}_\theta(\langle \mathbf{v}, \mathbf{X} \rangle) \geq \text{Var}_\theta(\langle \mathbf{v}, \mathbf{X} \rangle | \mathbf{X}_{S^c}) \quad (49)$$

$$= \sum_{i \in S} \text{Var}_\theta(v_i X_i | \mathbf{X}_{S^c}) \quad (50)$$

$$= \sum_{i \in S} v_i^2 \frac{1}{L_0(p)} \geq \frac{\|\mathbf{v}\|_2^2}{(k+1)L_0(p)}. \quad (51)$$

Here the identity in Eq. (50) follows because the $(X_i)_{i \in S}$ are conditionally independent given \mathbf{X}_{S^c} (note that S is an independent set in G , i.e. there is no edge connecting two vertices in S), and because $(X_i)_{i \in S^c}$ constant given \mathbf{X}_{S^c} . The expressions in Eq. (50) follow from Eqs. (47) and (48).

We therefore established also condition C2, with $L(p) = (k+1)L_0(p)$. This finishes the proof. \square

3.2 Tractable estimation from samples

In this section we assume to be given n i.i.d. samples $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n \sim p_\theta$ and denote by $\mathbf{X}^{(n)} = (\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n)$ the entire dataset. We will seek an estimator $\hat{\theta} = \hat{\theta}(\mathbf{X}^{(n)}; \xi)$ that can be computed efficiently, and is consistent in the sense that, for a sequence $\xi_n \rightarrow 0$

$$\hat{\theta}(\mathbf{X}^{(n)}; \xi_n) \xrightarrow{p} \theta, \quad (52)$$

in p_θ -probability.

It is indeed fairly easy to construct such an estimator: we will use an approach developed in [AKN06]. Fix $i \in [p]$ and say we want to estimate θ_i . Let $N_i(x_i; \mathbf{x}_{\partial i})$ be number of samples ℓ such that $X_i^{(\ell)} = x_i$, $\mathbf{X}_{\partial i}^{(\ell)} = \mathbf{x}_{\partial i}$. In formulae

$$N_i(x_i; \mathbf{x}_{\partial i}) = \#\{\ell \in [n] : X_i^{(\ell)} = x_i, \mathbf{X}_{\partial i}^{(\ell)} = \mathbf{x}_{\partial i}\}. \quad (53)$$

Then by the law of large numbers we have, almost surely and in probability

$$\lim_{n \rightarrow \infty} \frac{N_i(1; \mathbf{0})}{N_i(0; \mathbf{0})} = \frac{\mathbb{P}_\theta(X_i = 1, \mathbf{X}_{\partial i} = \mathbf{0})}{\mathbb{P}_\theta(X_i = 0, \mathbf{X}_{\partial i} = \mathbf{0})} = \frac{\mathbb{P}_\theta(X_i = 1 | \mathbf{X}_{\partial i} = \mathbf{0})}{\mathbb{P}_\theta(X_i = 0 | \mathbf{X}_{\partial i} = \mathbf{0})}. \quad (54)$$

Now, an immediate generalization of Eq. (43) yields

$$\frac{\mathbb{P}_\theta(X_i = 1 | \mathbf{X}_{\partial i} = \mathbf{0})}{\mathbb{P}_\theta(X_i = 0 | \mathbf{X}_{\partial i} = \mathbf{0})} = e^{2\beta k + \theta_i}. \quad (55)$$

This immediately suggests the estimator

$$\hat{\theta}_i \equiv -2\beta k + \log \frac{N_i(1; \mathbf{0})}{N_i(0; \mathbf{0})}. \quad (56)$$

This can be obviously evaluated for all vertices $i \in [p]$, in time linear in p and in n . We next provide a non-asymptotic estimate on the number of samples necessary to achieve a desired level of accuracy.

Proposition 3.2. *Let $\xi > 0$ be a precision parameter and Δ an error probability, and let $\boldsymbol{\theta}$ be such that $\|\boldsymbol{\theta}\|_\infty \leq \theta_{\max}$, $\theta_{\max} \geq 1$. Then, letting $\hat{\boldsymbol{\theta}}$ denote the estimator (56), we have*

$$\mathbb{P}_{\boldsymbol{\theta}}\left(\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_\infty \leq \xi\right) \geq 1 - \Delta, \quad (57)$$

provided $n \geq e^{C_* \theta_{\max}} \xi^{-2} \log(p/\Delta)$, where $C_* = C_*(k, \beta)$ is a constant.

Proof. Fix $i \in [p]$, and let, for $x \in \{0, 1\}$,

$$q_x \equiv \mathbb{P}_{\boldsymbol{\theta}}(X_i = x; \mathbf{X}_{\partial i} = \mathbf{0}). \quad (58)$$

Letting $D \subseteq [p]$, $|D| \leq k(k-1)$, be the set of neighbors of $\{i\} \cup \partial i$ in G , we have, by the Markov property

$$\min_{\mathbf{x}_D \in \{0,1\}^D} \mathbb{P}_{\boldsymbol{\theta}}(X_i = x; \mathbf{X}_{\partial i} = \mathbf{0} | \mathbf{X}_D = \mathbf{x}_D) \leq q_x \leq \max_{\mathbf{x}_D \in \{0,1\}^D} \mathbb{P}_{\boldsymbol{\theta}}(X_i = x; \mathbf{X}_{\partial i} = \mathbf{0} | \mathbf{X}_D = \mathbf{x}_D). \quad (59)$$

An explicit calculation shows that there exists a constant $C = C(k, \beta)$ such that

$$e^{-C\theta_{\max}} \leq q_x \leq 1 - e^{-C\theta_{\max}}. \quad (60)$$

Note that, for $x \in \{0, 1\}$, $N_i(x; \mathbf{0})$ is a binomial random variable $\text{Binom}(n, q_x)$. Standard tail bounds on binomial random variables yield (for $\varepsilon \leq 1/2$)

$$\mathbb{P}_{\boldsymbol{\theta}}\left(|N_i(x; \mathbf{0}) - \mathbb{E}N_i(x; \mathbf{0})| \geq \varepsilon \mathbb{E}N_i(x; \mathbf{0})\right) \leq 2 \exp\left\{-\frac{n\varepsilon^2}{8 \min(q_x, 1 - q_x)}\right\}. \quad (61)$$

Using Eq. (55) and the definition (56), together with Eq. (60), we get

$$\mathbb{P}_{\boldsymbol{\theta}}\left(|e^{\hat{\theta}_i} - e^{\theta_i}| \geq \varepsilon e^{\theta_i}\right) \leq 2 \exp\left\{-ne^{C\theta_{\max}}\varepsilon^2\right\}. \quad (62)$$

Passing to logarithms and using $|\log(1+x)| \leq C|x|$ for all $x \leq 1/2$, this yields, for all $\varepsilon < 1/4$, and eventually a different constant C ,

$$\mathbb{P}_{\boldsymbol{\theta}}\left(|\hat{\theta}_i - \theta_i| \geq \xi\right) \leq 2 \exp\left\{-ne^{C\theta_{\max}}\xi^2\right\}. \quad (63)$$

The proof is completed by choosing n so that the right hand side is upper bounded by Δ/p and using the union bound over the vertices p . \square

4 Discussion and related literature

From a mathematical point of view, the phenomenon highlighted by Theorem 1 is not new. It can be traced back to two well-understood facts, and one simple remark:

1. First well-understood fact. The log partition function is the value of a convex optimization problem:

$$\log Z(\mathbf{0}) = \max \left\{ F(\boldsymbol{\tau}) : \boldsymbol{\tau} \in [0, 1]^p \right\}. \quad (64)$$

2. Second well-understood fact. Recall that a weak evaluation oracle for F is an oracle that –on input $\boldsymbol{\tau}, \varepsilon$, with $\boldsymbol{\tau} \in [0, 1]^p$ and $\varepsilon > 0$ – returns \widehat{F} such that $|F(\boldsymbol{\tau}) - \widehat{F}| \leq \varepsilon$. Given such an oracle, then the optimization problem (64) can be solved in polynomial time with accuracy $C(p)\varepsilon$, with $C(p)$ a polynomial.

(See for instance [Lov87], Lemma 2.2.4 and Theorem 2.2.14 or [GLS81].)

3. Simple remark. We know that F is differentiable with Lipschitz continuous gradient for $\boldsymbol{\tau} \in [\delta, 1 - \delta]^p$, by assumption C2. Further, we showed quite easily that, letting $\boldsymbol{\tau}^{(0)} = (\delta, \dots, \delta)^\top$, we have $F(\boldsymbol{\tau}^{(0)}) \approx F(\mathbf{0}) = \log h(\mathbf{0})$. It follows that $F(\boldsymbol{\tau})$ can be –approximately– evaluated by ‘integrating’ the gradient ∇F between $\boldsymbol{\tau}^{(0)}$ and $\boldsymbol{\tau}$.

Since we assume to have access to a gradient oracle (i.e. an oracle for $\nabla F(\boldsymbol{\tau}) = -\boldsymbol{\theta}_*(\boldsymbol{\tau})$), we can construct an evaluation oracle. Hence, by the previous facts we can approximate $\log Z(\mathbf{0})$.

In other words, we could have proven Theorem 1 by using the construction at point 3, and then referring to standard results in convex optimization [GLS81, Lov87]. We preferred a self-contained proof, that uses additional structure of the problem (differentiability of F and existence of an oracle for ∇F).

The specific example discussed in Section 3 can be regarded as a graphical model. While we used a specific class of models, namely antiferromagnetic Ising models defined by Eq. (39), approximating the partition function of a graphical model is –in many cases– intractable [GJP03, Sly10, SS12, GSV12, CCGL12]. Hence, the conclusion is that –generally speaking– *it is intractable to estimate the parameters of a graphical model from sufficient statistics* (unless the model has special structure).

The literature on estimating parameters and structure of a graphical model is quite vast, and we can only provide a few pointers here. Traditional methods [HS83, AHS85] attempt at maximizing the likelihood function by gradient ascent. These approaches are necessarily based on sufficient statistics, and hence covered by our Theorem 1: in general, they cannot be implemented in polynomial time. The bottleneck is quite apparent in this case: evaluating the likelihood function or its gradient requires to compute expectations with respect to $p_{\boldsymbol{\theta}}$ which –in general– is hard. Notice that this difficulty is not circumvented by regularizing the likelihood function, as in the graphical LASSO [BEGd08].

A possible approach to overcome this problem was proposed in [RWL⁺10]: the basic idea is to reconstruct the neighborhood of a node i by performing a logistic regression against the other nodes. This approach generalizes to discrete graphical models a method developed in [MB06] for Gaussian graphical models. As shown in [BM09], the logistic regression method fails unless the graphical model satisfies a ‘weak dependence’ condition. Under the same type of condition, the log partition function can be computed efficiently.

Several papers [AKN06, CT06, BMS08, RSS12, ATHW12] develop algorithm to learn parameters and graph structures, with consistency guarantees under weak assumptions on the graphical model. As stressed in Section 3.2, these algorithms do not make use exclusively of the sufficient statistics and instead effectively estimate the joint distributions of subsets of k variables, with k depending on the maximum degree.

On the impossibility side, Bogdanov, Mossel and Vadhan [BMV08] showed that estimating graphical models with hidden nodes is intractable. Singh and Vishnoi [SV13] establish a result that is very similar to Theorem 1 with a slightly different oracle assumption. Their proof uses the ellipsoid algorithm instead of projected gradient, and their motivation is related to maximum entropy rounding techniques in optimization algorithms. Roughgarden and Kearns [RK13] establish equivalence of various computational tasks in graphical models. Again, they use convex duality and the ellipsoid algorithm.

Finally, I recently became aware of unpublished work by Bresler, Gamarnik and Shah [BGS14], that also proves intractability of learning graphical models from sufficient statistics. Their proof strategy is broadly similar to the one presented here but differs in the solution of several technical obstacles.

Acknowledgements

This work was partially supported by the NSF grant CCF-1319979 and the grants AFOSR/DARPA FA9550-12-1-0411 and FA9550-13-1-0036. I am grateful to Guy Bresler for discussing [BGS14] prior to publication.

A Simple properties of exponential families

In this section we provide a self-contained proof of Proposition 2.4. **Fact1** and **Fact2** are immediate and we omit their proof.

Fact3. For any vector $\mathbf{v} \in \mathbb{R}^p$, $\|\mathbf{v}\|_2 = 1$, we have

$$\langle \mathbf{v}, \text{Cov}_{\boldsymbol{\theta}}(\mathbf{X})\mathbf{v} \rangle = \text{Var}_{\boldsymbol{\theta}}(\langle \mathbf{v}, \mathbf{X} \rangle) = \frac{1}{2} \mathbb{E}_{\boldsymbol{\theta}} \left\{ \langle \mathbf{v}, \mathbf{X} - \mathbf{X}' \rangle^2 \right\}, \quad (65)$$

with $\mathbf{X}, \mathbf{X}' \sim p_{\boldsymbol{\theta}}$ independent. Letting $\mathbf{s} = \text{sign}(\mathbf{v})$, and continuing the above chain of inequalities, we get

$$\langle \mathbf{v}, \text{Cov}_{\boldsymbol{\theta}}(\mathbf{X})\mathbf{v} \rangle \geq 2 \cdot \frac{1}{2} p_{\boldsymbol{\theta}}(\mathbf{s}) p_{\boldsymbol{\theta}}(-\mathbf{s}) \langle \mathbf{v}, \mathbf{s} - (-\mathbf{s}) \rangle^2 \quad (66)$$

$$\geq p_{\boldsymbol{\theta}}(\mathbf{s}) p_{\boldsymbol{\theta}}(-\mathbf{s}) \|\mathbf{v}\|_1^2 \geq \min_{\mathbf{x} \in \{0,1\}^n} p_{\boldsymbol{\theta}}(\mathbf{x})^2, \quad (67)$$

which yields the desired claim.

Fact4. It is convenient to extend by continuity $\boldsymbol{\tau}_* : \overline{\mathbb{R}} \rightarrow [0, 1]^p$, where $\overline{\mathbb{R}} \equiv \mathbb{R} \cup \{+\infty, -\infty\}$ is the extended real line. It is a simple analysis exercise to check that this extension exists and is well defined. Indeed, if $\boldsymbol{\theta} \in \overline{\mathbb{R}}^n$ is such that $\theta_i = +\infty$ for all $i \in S_+ \subseteq [n]$, $\theta_i = -\infty$ for all $i \in S_- \subseteq [n]$,

and $\theta_i \in (-\infty, +\infty)$ for all $i \in S_0 \equiv [n] \setminus (S_+ \cup S_-)$, then we have

$$\tau_*(\theta) = \frac{1}{\tilde{Z}(\theta)} \sum_{\substack{\mathbf{x} \in \{0,1\}^n \text{ s.t.} \\ \mathbf{x}_{S_+}=1, \mathbf{x}_{S_-}=0}} \mathbf{x} h(\mathbf{x}) \exp \left\{ \langle \theta_{S_0}, \mathbf{x}_{S_0} \rangle \right\}, \quad (68)$$

$$\tilde{Z}(\theta) \equiv \sum_{\substack{\mathbf{x} \in \{0,1\}^n \text{ s.t.} \\ \mathbf{x}_{S_+}=1, \mathbf{x}_{S_-}=0}} h(\mathbf{x}) \exp \left\{ \langle \theta_{S_0}, \mathbf{x}_{S_0} \rangle \right\}, \quad (69)$$

We next prove that τ_* is injective. Indeed, assume by contradiction that $\tau_*(\theta_1) = \tau_*(\theta_2)$ for $\theta_1 \neq \theta_2$. Then, by the intermediate value theorem, there exists $\lambda \in [0, 1]$ such that, letting $\theta_\lambda = \lambda \theta_1 + (1 - \lambda) \theta_2$,

$$\tau_*(\theta_1) - \tau_*(\theta_2) = \nabla_{\theta} \tau_*(\theta_\lambda) (\theta_1 - \theta_2), \quad (70)$$

and hence, letting $\mathbf{v} = \theta_1 - \theta_2$,

$$0 = \langle \mathbf{v}, \tau_*(\theta_1) - \tau_*(\theta_2) \rangle = \langle \mathbf{v}, \text{Cov}_{\theta}(\theta_\lambda) \mathbf{v} \rangle, \quad (71)$$

which is impossible by Fact 3 above.

In order to prove that $\tau_* : \mathbb{R}^p \rightarrow [0, 1]^p$ is surjective, we will proceed by induction over the problem's dimension p . The claim is obvious for $p = 1$. We assume next that it holds for all dimensions up to $(p - 1)$ and prove it for dimension p . This claim follows by continuity, after proving that the image of τ_* contains the boundary $B_p \equiv [0, 1]^p \setminus (0, 1)^p$. Namely, for each $\tau \in B_p$, there exists $\theta \in \mathbb{R}^p$ such that $\tau_*(\theta) = \tau$.

To see that this is the case, fix one such τ . Let $S_+ \equiv \{i \in [n] : \tau_i = 1\}$, $S_- \equiv \{i \in [n] : \tau_i = 0\}$, and $S_0 \equiv [p] \setminus (S_+ \cup S_-)$ and note that, by assumption $|S_0| \leq n - 1$. Take θ such that $\theta_i = +\infty$ for all $i \in S_+$ and $\theta_i = -\infty$ for all $i \in S_-$. By Eq. (68), we have $\tau_*(\theta)_{S_+ \cup S_-} = \tau_{S_+ \cup S_-}$ and it is therefore sufficient to check the values of $\tau_*(\theta)$ on S_0 . Let $h_{S_+, S_-}(\mathbf{x}_{S_0}) = h(\mathbf{x})$ where $\mathbf{x}_{S_+} = 1_{S_+}$ and $\mathbf{x}_{S_-} = 0_{S_-}$. Then, again by Eq. (68), we have

$$\tau_*(\theta)_{S_0} = \frac{1}{\tilde{Z}(\theta)} \sum_{\mathbf{x}_{S_0} \in \{0,1\}^{S_0}} \mathbf{x}_{S_0} h_{S_+, S_-}(\mathbf{x}_{S_0}) \exp \left\{ \langle \theta_{S_0}, \mathbf{x}_{S_0} \rangle \right\}, \quad (72)$$

$$\tilde{Z}(\theta) = \sum_{\mathbf{x}_{S_0} \in \{0,1\}^{S_0}} h_{S_+, S_-}(\mathbf{x}_{S_0}) \exp \left\{ \langle \theta_{S_0}, \mathbf{x}_{S_0} \rangle \right\}. \quad (73)$$

Comparison with Eq. (11) shows that –by the induction hypothesis– there exists θ_{S_0} such that $\tau_*(\theta)_{S_0} = \tau_{S_0}$.

Fact5. This is a standard exercise with Legendre-Fenchel transforms and we omit its proof.

Fact6. Recall the Gibbs variational principle [MM09]

$$A(\theta) = \max_{q \in \mathcal{P}_p} \left\{ H(q) + \mathbb{E}_q(\log h(\mathbf{X}) + \langle \theta, \mathbf{X} \rangle) \right\}. \quad (74)$$

The claim follows by comparing this expression with Eq. (21).

B Finishing the proof of Theorem 1

In this appendix we complete the proof of Theorem 1. In the simplified version presented in Section 2, we assumed to have access to an oracle returning $\theta_*(\tau)$ when queried with value τ . We will show that our claim remains correct if the oracle is replaced by a polynomial consistent estimator $\hat{\theta}(\cdot, \cdot)$. By Definition 2.1, for any $\xi > 0$ we have

$$\|\hat{\theta}(\tau, \xi) - \theta_*(\tau)\|_2 \leq \xi. \quad (75)$$

The oracle $\theta_*(\cdot)$ is used in two points in the proof of Theorem 1:

1. In the implementation of the projected gradient algorithm, cf. Eq. (29). It is queried a total of $t_0(p, \varepsilon) \equiv 2pL(p)/\varepsilon$ times for this purpose.
2. In calculating the approximation of $Z(\mathbf{0})$, cf. Eq. (35). It is queried a total of $m_0(p, \varepsilon) \equiv \lceil 4L(p)p/\varepsilon \rceil$ times for this purpose.

We will replace these queries with queries to $\hat{\theta}(\cdot, \xi)$ with $1/\xi$ polynomial in p and $1/\varepsilon$. Since the total number of calls is also polynomial in p and $1/\xi$, this yields an algorithm that is polynomial in p and $1/\varepsilon$.

Queries for projected gradient. We denote by $\sigma^0, \sigma^1, \dots, \sigma^{t_0}$ the sequence generated by the projected gradient algorithm, with $\nabla_{\tau} F(\tau)$ approximated by $-\hat{\theta}(\tau, \xi)$. Namely, we have $\sigma^0 = \tau^0$ and, for all $t \geq 0$,

$$\sigma^{t+1} = P_{\delta} \left(\sigma^t - \frac{1}{L} \hat{\theta}(\sigma^t, \xi) \right). \quad (76)$$

Comparing with Eq. (29), we have (dropping the argument ξ of $\hat{\theta}$ in order to simplify the notation):

$$\|\tau^{t+1} - \sigma^{t+1}\|_2 \leq \left\| \left(\tau^t - \frac{1}{L} \hat{\theta}(\tau^t) \right) - \left(\sigma^t - \frac{1}{L} \hat{\theta}(\sigma^t) \right) \right\|_2 \quad (77)$$

$$\leq \left\| \tau^t + \frac{1}{L} \nabla_{\tau} F(\tau^t) - \sigma^t - \frac{1}{L} \nabla_{\tau} F(\sigma^t) \right\|_2 \quad (78)$$

$$+ \frac{1}{L} \|\hat{\theta}(\sigma^t) - \theta_*(\sigma^t)\|_2 + \frac{1}{L} \|\hat{\theta}(\tau^t) - \theta_*(\tau^t)\|_2 \quad (79)$$

$$\leq \left\| \left(I + \frac{1}{L} \nabla^2 F(\rho^t) \right) (\tau^t - \sigma^t) \right\|_2 + \frac{2\xi}{L}, \quad (80)$$

where the last inequality follows from the definition of $\hat{\theta}$, and by applying the intermediate value theorem, with ρ^t a point on the interval $[\tau^t, \sigma^t]$. Now, by Eq. (20), (23) and assumption C2, we have $-L \preceq \nabla^2 F(\rho^t) \preceq \mathbf{0}$, and therefore $\|I + L^{-1} \nabla^2 F(\rho^t)\|_2 \leq 1$ (in operator norm). We get therefore

$$\|\tau^{t+1} - \sigma^{t+1}\|_2 \leq \|\tau^t - \sigma^t\|_2 + \frac{2\xi}{L}, \quad (81)$$

which, of course, implies

$$\|\tau^{t_0} - \sigma^{t_0}\|_2 \leq \frac{2t_0 \xi}{L}. \quad (82)$$

Notice that, for any $\boldsymbol{\tau} \in D_p(\delta)$,

$$\|\nabla_{\boldsymbol{\tau}} F(\boldsymbol{\tau})\|_2 = \|\nabla_{\boldsymbol{\tau}} F(\boldsymbol{\tau}) - \nabla_{\boldsymbol{\tau}} F(\boldsymbol{\tau}_*)\|_2 \leq L \|\boldsymbol{\tau} - \boldsymbol{\tau}_*\|_2 \leq L\sqrt{p}. \quad (83)$$

Hence

$$|F(\boldsymbol{\tau}^{t_0}) - F(\boldsymbol{\sigma}^{t_0})| \leq L\sqrt{p} \|\boldsymbol{\tau}^{t_0} - \boldsymbol{\sigma}^{t_0}\|_2 \leq 2t_0 \sqrt{p}\xi. \quad (84)$$

In particular, by choosing $\xi \leq \varepsilon/(16t_0(p, \varepsilon)\sqrt{p})$, it follows that $|F(\boldsymbol{\tau}^{t_0}) - F(\boldsymbol{\sigma}^{t_0})| \leq \varepsilon/8$.

Queries for evaluating Eq. (35). We let $\boldsymbol{\sigma}^{(0)} = \boldsymbol{\tau}^{(0)}$, $\boldsymbol{\sigma}^{(m)} = \boldsymbol{\sigma}^{t_0}$ and $\boldsymbol{\sigma}^{(\ell)}$, $\ell \in \{1, \dots, m-1\}$ be given by interpolating linearly. The approximation \widehat{Z} in Eq. (35) is replaced by

$$\log \widehat{Z} = \log h(\mathbf{0}) - \sum_{\ell=1}^m \langle \boldsymbol{\theta}_*(\boldsymbol{\tau}^{(\ell)}), \boldsymbol{\tau}^{(\ell)} - \boldsymbol{\tau}^{(\ell-1)} \rangle. \quad (85)$$

The approximation error is bounded analogously to Eq. (36). We get two additional error terms

$$\left| \log \frac{\widehat{Z}}{Z} \right| \leq \frac{3\varepsilon}{4} + |F(\boldsymbol{\tau}^{t_0}) - F(\boldsymbol{\sigma}^{t_0})| + \left| \sum_{\ell=1}^m \langle \widehat{\boldsymbol{\theta}}(\boldsymbol{\sigma}^{(\ell)}) - \boldsymbol{\theta}_*(\boldsymbol{\sigma}^{(\ell)}), \boldsymbol{\sigma}^{(\ell)} - \boldsymbol{\sigma}^{(\ell-1)} \rangle \right| \quad (86)$$

$$\leq \frac{7\varepsilon}{8} + \|\boldsymbol{\sigma}^{(m)} - \boldsymbol{\sigma}^{(0)}\|_2 \max_{\ell \in [m]} \|\widehat{\boldsymbol{\theta}}(\boldsymbol{\sigma}^{(\ell)}) - \boldsymbol{\theta}_*(\boldsymbol{\sigma}^{(\ell)})\|_2 \quad (87)$$

$$\leq \frac{7\varepsilon}{8} + \xi\sqrt{p}. \quad (88)$$

The latter is bounded by ε as soon as $\xi \leq 1/(8\varepsilon\sqrt{p})$, which is guaranteed since we already required $\xi \leq \varepsilon/(16t_0(p, \varepsilon)\sqrt{p})$.

We conclude that the desired approximation error is guaranteed by using the oracle $\widehat{\boldsymbol{\theta}}$, with accuracy parameter $\xi \leq \varepsilon/(16t_0(p, \varepsilon)\sqrt{p})$. This concludes the proof, since it corresponds to $(1/\xi)$ polynomial in p and $(1/\varepsilon)$.

References

- [AB09] Sanjeev Arora and Boaz Barak, *Computational complexity: a modern approach*, Cambridge University Press, 2009.
- [AHS85] David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski, *A learning algorithm for boltzmann machines*, Cognitive science **9** (1985), no. 1, 147–169.
- [AKN06] Pieter Abbeel, Daphne Koller, and Andrew Y Ng, *Learning factor graphs in polynomial time and sample complexity*, The Journal of Machine Learning Research **7** (2006), 1743–1788.
- [ATHW12] Animashree Anandkumar, Vincent YF Tan, Furong Huang, and Alan S Willsky, *High-dimensional structure estimation in Ising models: Local separation criterion*, The Annals of Statistics **40** (2012), no. 3, 1346–1375.

- [BEGd08] Onureena Banerjee, Laurent El Ghaoui, and Alexandre d’Aspremont, *Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data*, The Journal of Machine Learning Research **9** (2008), 485–516.
- [BGS14] Guy Bresler, David Gamarnik, and Devavrat Shah, *Hardness of parameter estimation in graphical models*, unpublished, 2014.
- [BM09] José Bento and Andrea Montanari, *Which graphical models are difficult to learn?*, Neural Information Processing Systems (Vancouver), December 2009.
- [BMS08] Guy Bresler, Elchanan Mossel, and Allan Sly, *Reconstruction of markov random fields from samples: Some observations and algorithms*, Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques, Springer, 2008, pp. 343–356.
- [BMV08] Andrej Bogdanov, Elchanan Mossel, and Salil Vadhan, *The complexity of distinguishing markov random fields*, Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques, Springer, 2008, pp. 331–342.
- [BT09] Amir Beck and Marc Teboulle, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM Journal on Imaging Sciences **2** (2009), no. 1, 183–202.
- [CCGL12] Jin-Yi Cai, Xi Chen, Heng Guo, and Pinyan Lu, *Inapproximability after uniqueness phase transition in two-spin systems*, Combinatorial Optimization and Applications, Springer, 2012, pp. 336–347.
- [CT06] Imre Csiszár and Zsolt Talata, *Consistent estimation of the basic neighborhood of Markov random fields*, The Annals of Statistics (2006), 123–145.
- [Efr78] Bradley Efron, *The geometry of exponential families*, The Annals of Statistics **6** (1978), no. 2, 362–376.
- [GJP03] Leslie Ann Goldberg, Mark Jerrum, and Mike Paterson, *The computational complexity of two-state spin systems*, Random Structures & Algorithms **23** (2003), no. 2, 133–154.
- [GLS81] Martin Grötschel, László Lovász, and Alexander Schrijver, *The ellipsoid method and its consequences in combinatorial optimization*, Combinatorica **1** (1981), no. 2, 169–197.
- [Gri06] Geoffrey R Grimmett, *The random-cluster model*, vol. 333, Springer Science & Business Media, 2006.
- [GSV12] Andreas Galanis, Daniel Stefankovic, and Eric Vigoda, *Inapproximability of the partition function for the antiferromagnetic ising and hard-core models*, arXiv:1203.2226 (2012).
- [Hop82] John J Hopfield, *Neural networks and physical systems with emergent collective computational abilities*, Proceedings of the National Academy of Sciences **79** (1982), no. 8, 2554–2558.
- [HS83] GE Hinton and TJ Sejnowski, *Analysing cooperative computation*, Proceedings of the Fifth Annual Conference of the Cognitive Science Society, 1983.

- [Isi25] Ernst Ising, *Beitrag zur theorie des ferromagnetismus*, Zeitschrift für Physik A Hadrons and Nuclei **31** (1925), no. 1, 253–258.
- [LC98] EL Lehmann and George Casella, *Theory of point estimation*, 2 ed., Springer, 1998.
- [Lov87] László Lovász, *An algorithmic theory of numbers, graphs and convexity*, vol. 50, SIAM, 1987.
- [MB06] Nicolai Meinshausen and Peter Bühlmann, *High-dimensional graphs and variable selection with the lasso*, Ann. Statist. **34** (2006), 1436–1462.
- [MM09] Marc Mézard and Andrea Montanari, *Information, Physics and Computation*, Oxford, 2009.
- [RK13] Tim Roughgarden and Michael Kearns, *Marginals-to-models reducibility*, Advances in Neural Information Processing Systems, 2013, pp. 1043–1051.
- [RSS12] Avik Ray, Sujay Sanghavi, and Sanjay Shakkottai, *Greedy learning of graphical models with small girth*, Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on, IEEE, 2012, pp. 2024–2031.
- [RWL⁺10] Pradeep Ravikumar, Martin J Wainwright, John D Lafferty, et al., *High-dimensional Ising model selection using ℓ_1 -regularized logistic regression*, The Annals of Statistics **38** (2010), no. 3, 1287–1319.
- [Sly10] Allan Sly, *Computational transition at the uniqueness threshold*, Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on, IEEE, 2010, pp. 287–296.
- [SS12] Allan Sly and Nike Sun, *The computational hardness of counting in two-spin models on d -regular graphs*, Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on, IEEE, 2012, pp. 361–369.
- [SST14] Alistair Sinclair, Piyush Srivastava, and Marc Thurley, *Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs*, Journal of Statistical Physics **155** (2014), no. 4, 666–686.
- [SV13] Mohit Singh and Nisheeth K Vishnoi, *Entropy, optimization and counting*, arXiv preprint arXiv:1304.8108 (2013).